



Cisco Nexus Dashboard Infrastructure Management, Release 3.1.x

Table of Contents

System Settings	1
Persistent IP Addresses	5
Persistent IP Guidelines and Limitations	5
Enabling BGP On All Nodes	6
Configuring Persistent IPs	7
Multi-Cluster Connectivity	9
Guidelines and Limitations	9
Connecting Multiple Clusters	10
Disconnecting Clusters	12
Managing Secondary Nodes	13
Adding Secondary Nodes	13
Deleting a Secondary Node	14
Managing Standby Nodes	15
Adding Standby Nodes	15
Replacing Single Primary Node with Standby Node	16
Replacing Two Primary Nodes with Standby Nodes	17
Deleting Standby Nodes	22
Trademarks	23

System Settings

The **System Settings** GUI screen allows you to configure a number of options specific to the Nexus Dashboard cluster and its nodes. It will also display information about any issues that may be present in your Nexus Dashboard cluster.

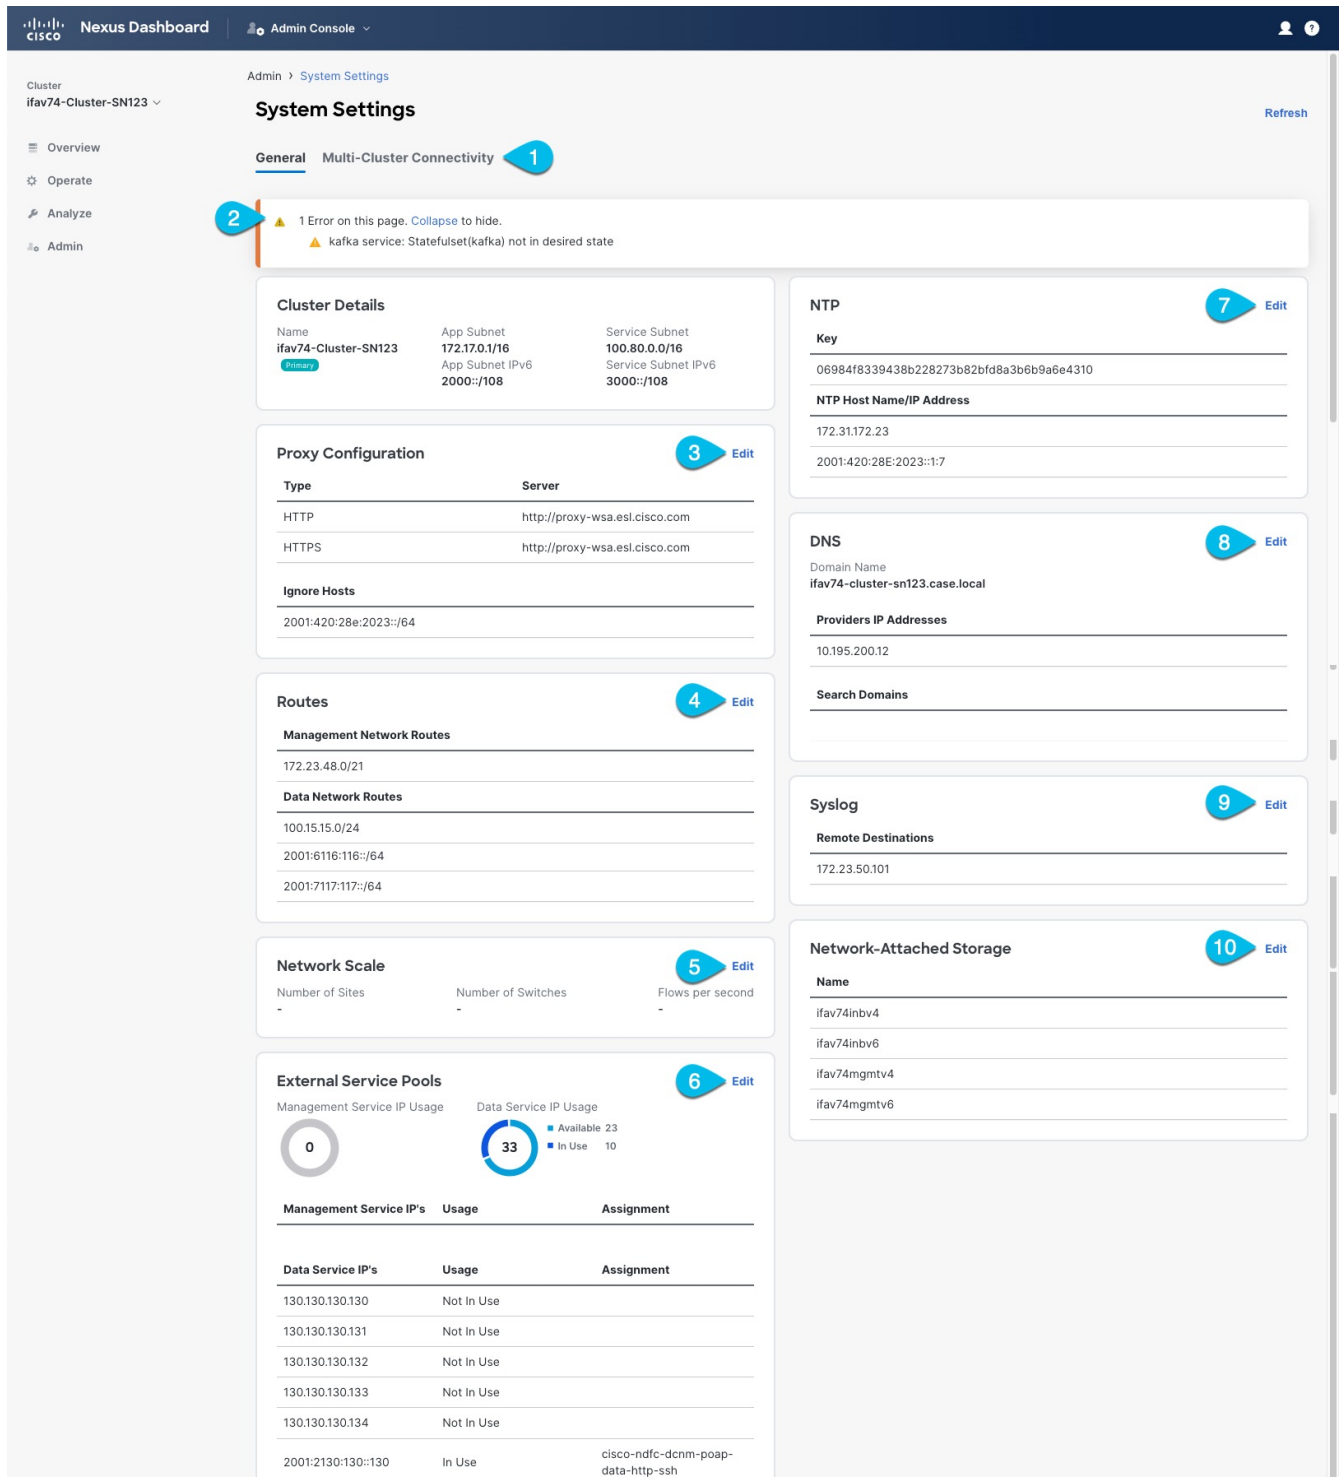


Figure 1. System Settings

1. The **Multi-cluster Connectivity** tab allows you to connect multiple clusters together for a single pane of glass view and administration of the clusters and their sites, services, and configurations.

For more information, see [Multi-Cluster Connectivity](#).

2. The errors and warning tile will display the number of existing issues in your cluster. You can click **Expand** to see the full list of specific issues.
3. To configure a proxy for the Nexus Dashboard, click the **Edit** icon in the **Proxy Configuration** tile.

In certain deployment scenarios, such as with a combination of on-premises and cloud sites and the Nexus Dashboard cluster deployed inside a corporate network, you may have to access the Internet and the cloud sites through a proxy.



This release supports adding a single proxy server.

Note that Nexus Dashboard uses 2 main route tables—one for the Management network and one for the Data network—and by default, it will use the routing table of the originating IP address. In other words, Nexus Dashboard will attempt to reach the proxy from the routing table of the POD/Service that is trying to use the proxy.

For example, if you configure a proxy and establish Intersight connectivity from your Nexus Dashboard and then attempt to configure AppD integration from the Insights service running in the cluster, you may get an error stating that AppD host is not reachable. This happens because the proxy is only accessible from the management interface, so in such cases you also need to add a management network route for the proxy IP address, as described in "Management Network or Data Network routes" below.

To add a proxy server:

- a. Click **+Add Server** in the proxy configuration window.
- b. From the **Type** dropdown, select the type of traffic that you want to be proxied.
- c. In the **Server** field, provide the full address for the proxy server including the port if required.

For example <http://proxy.company.com:80>.

- d. If the server requires login credentials, provide the **Username** and **Password**.
- e. (Optional) Click **Add Ignore Host** to provide any hosts that will ignore the proxy.

You can add one or more hosts with which the cluster will communicate directly bypassing the proxy.

4. To add one or more Management Network or Data Network routes, click the **Edit** icon in the **Routes** tile.

Here you can define static routes for the management or data interfaces. For example, adding **10.195.216.0/21** as a Data Network route will cause all traffic destined to that subnet to transit out of the data network interface.

- o To add a management network route, click **Add Management Network Routes** and provide the destination subnet.
- o To add a data network route, click **Add Data Network Routes** and provide the destination subnet.

5. To add one or more External Service Pools, click the **Edit** icon in the **External Service Pools** tile.

This allows you to provide persistent IP addresses for services that require to retain the same IP

addresses even in case they are relocated to a different Nexus Dashboard node.

For detailed information and configuration steps, see [Persistent IP Addresses](#).

6. To configure NTP settings, click the **Edit** icon in the **NTP** tile.

By default, the NTP server that you configured when deploying the Nexus Dashboard cluster is listed here.

You can provide additional NTP servers by clicking **+Add NTP Server**.

You can remove existing NTP server by clicking the **Delete** icon next to it. Keep in mind that at least one NTP server must be configured in your cluster.

7. To configure DNS settings, click the **Edit** icon in the **DNS** tile.

By default, the DNS server and search domain that you configured when deploying the Nexus Dashboard cluster are listed here.

You can provide additional DNS servers and search domains by clicking **+Add a Provider** or **+Add a Search Domain** respectively.

You can remove existing DNS server by clicking the **Delete** icon next to it.

8. To provide one or more **syslog** servers to stream event logs to, click the **Edit** icon in the **Syslog** tile.

In the **Syslog** dialog that opens, click **+Add Remote Destinations** to add a new server. Then provide the IP address, protocol, and port number for the server and choose whether you want to enable streaming to this syslog server at this time.

For more information, see [History and Logs](#).

9. To configure Network-Attached Storage (NAS), click the **Edit** icon in the **Network-Attached Storage** tile.

Beginning with release 3.0(1), you can add a NAS server at the Nexus Dashboard level which can be utilized by the services running in your cluster.



This release supports only NFSv3 remote storage.

To add a NAS:

- a. Click **+Add Network-Attached Storage** in the NAS configuration window.
- b. Choose whether Nexus Dashboard has **Read Only** or **Read Write** access to this server.
- c. Provide the **Name** for the NAS server.
- d. (Optional) Provide the **Description**.
- e. Provide the **IP address** used to connect to the server.
- f. Specify the **Port** used to establish the connection if it is different from the default port **2049**.
- g. Provide the **Export Path** to a directory on the NAS server where information will be stored or read.

h. Specify the **Alert Threshold**.

i. Specify the storage **Limit**.

This limits the amount of storage that can be requested on the server by Nexus Dashboard. You can provide the values in Mibibytes or Gibibytes, for example **300Mi** or **10Gi**.

j. From the **Allowed Apps** dropdown, select which Nexus Dashboard services can access this storage.

Persistent IP Addresses

You can provide persistent IP addresses for services that require to retain the same IP addresses even in case they are relocated to a different Nexus Dashboard node.

Nexus Dashboard Insights requires some services (such as SNMP trap, syslog, and others) to stream data from the switches in your fabrics to the service. An IP address is configured on the switches for this purpose. Typically, if the IP address changes when the service is relocated, the service will reconfigure the new IP address on the switches.

In order to avoid this IP reconfiguration impact on the fabric switches, the service can request that the services IP addresses are preserved, in which case you will need to define a set of IP addresses which can be assigned to the service for this purpose.

If a service requires persistent IP addresses, you will not be able to enable that service in the Nexus Dashboard until enough IP addresses are defined as described below.



This feature is supported for Nexus Dashboard Insights with NDFC fabrics only. In addition, if you are using Layer 2 functionality only (IPs configured as part of the management and data subnets) and your Nexus Dashboard is deployed in VMware ESX, you must enable promiscuous mode for both management and data network interface portgroups, as described in <https://kb.vmware.com/s/article/1004099>.

Prior to Release 2.2(1), this feature was supported only for clusters where all nodes were part of the same Layer 3 network and the persistent IPs were defined as part of the node's management or data networks. Here the application uses Layer 2 mechanisms like Gratuitous ARP or Neighbor Discovery to advertise the persistent IPs within its Layer 3 network.

Beginning with Release 2.2(1), the feature is supported even if you deploy the cluster nodes in different Layer 3 networks. In this case, the persistent IPs are advertised out of each node's data links via BGP, which we refer to as "Layer 3 mode". The IPs must not overlap with any of the nodes' management or data subnets. If the persistent IPs are outside the data and management networks, this feature will operate in Layer 3 mode by default; if the IPs are part of those networks, the feature will operate in Layer 2 mode.

Persistent IP Guidelines and Limitations

When configuring persistent IPs for your services:

- Ensure that you check the documentation for the services you plan to deploy as some services do not support this feature or require additional guidelines.

At this time, Persistent IPs are supported for Nexus Dashboard Insights and Nexus Dashboard Fabric Controller. You can find the service-specific documentation at the following links:

- [Nexus Dashboard Fabric Controller](#)
- [Nexus Dashboard Insights](#)

- You can choose which mode you want to operate in as long as the following conditions apply:
 - If you choose to operate in Layer 2 mode, the nodes must be part of the same data and management networks.

- o If you choose to operate in Layer 3 mode, all nodes must have BGP configuration provided either during cluster deployment or after as described in [Enabling BGP On All Nodes](#).
- o You can switch between the two modes, in which case the existing services of a particular mode must be completely deleted and you will need to configure new persistent IPs corresponding to the new mode.
- If you configure one or more persistent IPs in Layer 3 mode and back up cluster configuration, the BGP settings required for this feature are not saved in the backup.

As such, you must ensure that you configure BGP for all cluster nodes before restoring any cluster configuration that contains Layer 3 persistent IPs in that cluster. If BGP is not configured prior to the configuration import, the import will fail.

Enabling BGP On All Nodes

If you want to operate in Layer 3 mode, you must enable and configure BGP for all nodes in your cluster. If you already configured BGP for each node during cluster deployment or if you want to operate in Layer 2 mode instead, you can skip this section and simply provide one or more persistent IPs from the nodes' management and data subnets, as described in [Configuring Persistent IPs](#). Note that if you choose to operate in Layer 2 mode, all nodes must be part of the same Layer 3 network. If you choose to operate in Layer 3 mode, at least one BGP peer must be configured on all cluster nodes to advertise the IPv4 or the IPv6 persistent IP addresses as described in this section.

Before you begin

- Ensure that the uplink peer routers are capable of exchanging the advertised persistent IPs across the Layer 3 networks of the cluster nodes.
- When a service requests a persistent IP address, the route advertised from the data links via BGP on the node where the service is running is maintained throughout the lifecycle of the service.

To configure BGP on the nodes:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the left navigation menu, select **System Resources > Nodes**.
3. Click the **Actions (...)** menu next to one of the nodes and choose **Edit**.
4. In the **Edit Node** screen, turn on **Enable BGP**.
5. In the **ASN** field, provide the autonomous system number for the node.
6. Click **+Add IPv4 BGP Peer** or **+Add IPv6 BGP Peer** to provide peer IP address information.
 - a. In the **Peer Address** field, provide the IPv4 or IPv6 address of the peer router for this node.

Multi-hop BGP peering is not supported, so you must ensure that the **Peer Address** is part of the node's data subnet.

- b. In the **Peer ASN** field, provide the autonomous system number of the peer router.

Only EBGP is supported, so you must ensure that the node ASN and Peer ASN are different.

- c. Click **Save** to save the changes.

7. Repeat these steps for every node in the cluster.

Every node in the cluster must have BGP configured.

You can configure the same ASN for all nodes or a different ASN per node

Configuring Persistent IPs

Before you begin

- For all persistent IPs, you must use either the Layer 2 or Layer 3 approach; a combination of the two is not supported.

If all nodes are in the same Layer 3 network, you can choose to use either the Layer 2 mode or Layer 3 mode for this feature. The two modes are described in [Persistent IP Addresses](#).

If the nodes are in different Layer 3 networks, you must configure the persistent IPs such that they don't overlap with either the management or the data subnets of the nodes.

- If the nodes in your cluster belong to different Layer 3 networks, you must have BGP enabled and configured as described in [Enabling BGP On All Nodes](#).
- There may be a momentary traffic interruption while a service using a persistent IP is relocated to a different node.

The interruption duration depends on the following factors:

- Time to detect the node failure
- Time for the service to get rescheduled to a different node
- Time for the service's external IP to get advertised from the scheduled node via GARP (IPv4) or neighbor discovery (IPv6) in case of Layer 2 mode
- Time for the service's external IP to get advertised from the scheduled node via BGP in case of layer 3 mode

To provide one or more persistent IP addresses:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the left navigation menu, select **Admin > System Settings**.
3. In the **External Service Pools** tile, click the **Edit** icon.
4. In the **External Service Pools** screen that opens, click **+Add IP Address** to add one or more IP addresses for the management or data networks.

When editing persistent IPs, the following rules apply:

- If all nodes in your cluster are part of the same Layer 3 network, you can choose one of the following:
 - Layer 2 mode, in which case the IP addresses you add for management services must be part of the management subnet and the IP addresses for data services must be part of the data subnet.
 - Layer 3 mode, in which case the IP addresses you add must not overlap with the management or the data subnets of the nodes. In this case, adding IPs under "Management Service IPs" is not supported and you must add the IPs to the "Data

Service IPs" category in the GUI.

- o You must provide either IPv4 or IPv6 IP addresses, you cannot give both.
- o You must add individual IP addresses one by one without any prefix; adding a range of IP addresses is not supported.
- o You can remove any previously defined IPs, but you will not be able to remove any IPs that are currently in use by one or more services.

Multi-Cluster Connectivity

You can establish connectivity between multiple Nexus Dashboard clusters for ease of access to all the clusters, as well as access to any of the sites and services running on any of the connected clusters.

When you add a second cluster, a group of clusters is formed. The cluster from which you create the group becomes the "primary" cluster with a number of unique characteristics that do not apply to other clusters in the group:

- You must use the primary cluster to connect all additional clusters.
- You must use the primary cluster to remove any of the clusters from the group.
- When upgrading Nexus Dashboard, you must upgrade the primary cluster before any other clusters in the group.

Establishing multi-cluster connectivity does not create any single databases with information from all clusters in the group. Every cluster continues to maintain its own configuration databases, while simultaneously being able to function as a proxy for all other clusters in the group regardless of which cluster an action or request is originated from or destined to.

Guidelines and Limitations

The following guidelines apply when configuring multi-cluster connectivity:

- This release supports multi-cluster connectivity between clusters deployed using physical or virtual (ESX) form factors only.

In other words, you can join physical Nexus Dashboard clusters with virtual (ESX) clusters, but virtual (KVM) or cloud clusters do not support this feature.

- For supported scale limits, such as number of clusters that can be connected together and number of sites across all clusters, see the [Nexus Dashboard Release Notes](#) for your release.
- Connectivity (HTTPS) must be established between the management interfaces of all the nodes of all the clusters, which will be connected via multi-cluster connectivity.
- The names of the sites onboarded in the clusters that you plan to connect together must be unique across those clusters.

Duplicate site names across different clusters may result in DNS resolution failures.

- The primary cluster, which you use to establish multi-cluster connectivity, must be running the same or a later release of Nexus Dashboard than any other cluster in the group.

In other words, you cannot connect a Nexus Dashboard cluster running release 2.3.1 from a primary cluster that is running release 3.0.1.

- If you are upgrading multiple clusters that are connected together, you must upgrade the primary cluster first.
- From any cluster in the connected clusters group, you can view other clusters only if they are running the same or earlier version of Nexus Dashboard.

In other words, if **cluster1** is running release 2.3.1 and **cluster2** is running release 2.2.1, you can view **cluster2** from **cluster1** but not vice versa.

- Multi-Cluster connectivity is supported for remote users only.

If you connect multiple clusters, but then login to one of the clusters as a local **admin** user, you will only be able to view and manage the local cluster into which you logged in.

To view and manage all clusters in the group, you must login as a remote user that is configured on all clusters.

- Nexus Dashboard Insights service in each cluster can view site groups from other Insights services across any cluster in the group.

However, when creating site groups, each Insights service can add sites which are onboarded in the same cluster where the service is installed only.

- Nexus Dashboard Orchestrator service supports managing only sites which are onboarded in the same cluster where the service is installed.

Connecting Multiple Clusters

Before you begin

- You must have familiarized yourself with the information provided in the [Guidelines and Limitations](#) section.
- You must have set up remote authentication and users on all clusters which you plan to connect.

Multi-Cluster connectivity is supported for remote users only, so you must configure the same remote user with **admin** privileges for all clusters. For additional details, see [Remote Authentication](#).

To connect another cluster:

1. Log in to the Nexus Dashboard GUI of the cluster which you want to designate as the primary.
2. Add second cluster.
 - a. From the main navigation menu, select **Admin > System Settings**.
 - b. In the main pane, select the **Multi-Cluster Connectivity** tab.
 - c. Click **Connect Cluster**.
3. Provide cluster information.
 - a. In the information fields, provide the hostname or IP address and the authentication information for the cluster you are adding.

You only need to provide the management IP address of one of the nodes in the target cluster. Other nodes' information will be automatically synced after connectivity is established.

- b. Then click **Save**.

The user you provide must have administrative rights on the cluster you are adding. The user credentials are used once when you are first establishing connectivity to the additional cluster. After initial connectivity is established, all subsequent communication is done through secure

keys. The secure keys are provisioned to each cluster while adding it to the group.

The cluster you are adding must not be part of an already existing group of clusters.

4. Repeat the procedure for any additional Nexus Dashboard cluster which you want to add to the group.

After multiple clusters are added to the group, you can see their status in the **Cluster Configuration > Multi-Cluster Connectivity** page.

Note that while you can view and manage any cluster from any other cluster as long as they are part of the same multi-cluster group, you can only add and remove clusters from the group when viewing the **primary** cluster.

The **Multi-Cluster Connectivity** page displays all clusters that are part of the multi-cluster group. The **Connect Cluster** button is shown only when viewing the primary cluster. To modify the cluster group, you need to navigate to the primary cluster, at which point the **Connect Cluster** button becomes available:

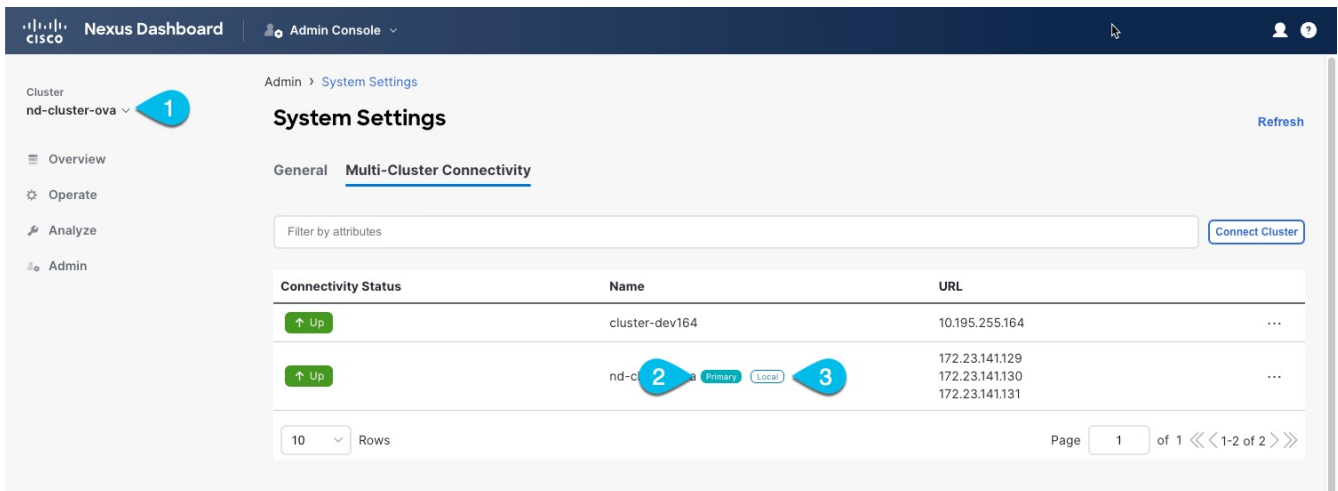


Figure 2. Primary vs Non-primary Clusters

1. The **Cluster: <name>** dropdown in the main navigation menu shows the cluster you are currently viewing.

You can select a different cluster from this dropdown, which opens a new window allowing you to navigate to another cluster in the same group.



While the 2.x releases of Nexus Dashboard allowed you to view and manage any cluster from any other cluster as long as they were part of the same multi-cluster group, release 3.0.1 changed this behavior. You can now easily navigate between clusters by picking a specific cluster from the **Cluster** dropdown in the main navigation pane, but you cannot manage or configure another cluster directly from the one where you are logged in.

2. The **Primary** label indicates the group's primary cluster.

You must be viewing this cluster to make any changes to the cluster group, such as adding or removing clusters.

3. The **Local** label indicates the cluster you logged into.

This is the cluster whose address is displayed in the browser's URL field. If you navigate to a different cluster as mentioned above, the browser URL and the **Local** label will not change.

Disconnecting Clusters

To disconnect a cluster from an existing group:

1. Log in to the Nexus Dashboard GUI of the primary cluster.

Adding and removing clusters from the group must be done from the primary cluster.

2. From the main navigation menu, select **Admin > System Settings**.
3. In the main pane, select the **Multi-Cluster Connectivity** tab.
4. From the **Actions (...)** menu for the cluster you want to remove, select **Disconnect Cluster**
5. In the confirmation window, click **Ok**.

Managing Secondary Nodes

You can add a number of secondary nodes to an existing 3-node cluster for horizontal scaling to enable application co-hosting.

For additional information about application co-hosting and cluster sizing, see the [Platform Overview](#) section of this document.



Secondary nodes are not supported for cloud form factors of Nexus Dashboard clusters deployed in AWS or Azure.

Adding Secondary Nodes

This section describes how to add a secondary node to your cluster to enable horizontal scaling

Before you begin

- Ensure that the existing primary nodes and the cluster are healthy.
- Prepare and deploy the new node.
- Ensure that the node you are adding is powered on.
- If you are adding a physical node, ensure that you have the new node's CIMC IP address and login information.

You will need to use the CIMC information to add the new node using the Nexus Dashboard GUI.

- If you are adding a virtual node, ensure that you have the node's management IP address and login information.

To add a secondary node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. In the main pane, click **Add Node**.

The **Add Node** screen opens.

4. In the **Add Node** screen, provide the node information.
 - a. Provide the **Name** of the node.
 - b. From the **Type** dropdown, select **secondary**.
 - c. Provide the **Credentials** information for the node, then click **Verify**.

For physical nodes, this is the IP address, username, and password of the server's CIMC. The CIMC will be used to configure the rest of the information on the node.

For virtual nodes, this is the IP address and **rescue-user** password you defined for the node when deploying it.

- d. Provide the **Management Network** information.

For virtual nodes, the management network information will be pre-populated with the information pulled from the node based on the IP address and credentials you provided in the previous sub-step.

For physical nodes, you must provide the management network IP address, netmask, and gateway now.

- e. Provide the **Data Network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

- f. (Optional) Provide IPv6 information for the management and data networks.

Starting with release 2.1.1, Nexus Dashboard supports dual stack IPv4/IPv6 for the management and data networks.

If you want to provide IPv6 information, you must do it when adding the node.

All nodes in the cluster must be configured with either only IPv4 or dual IPv4/IPv6 stack.

5. Click **Save** to add the node.

The configuration will be pushed to the node and the node will be added to the list in the GUI.

6. If you are running Nexus Dashboard Insights and you added a secondary node to replace a faulty secondary node (RMA), disable and re-enable Nexus Dashboard Insights.

Nexus Dashboard Insights must be restarted to properly redistribute the services to the new secondary node.

Deleting a Secondary Node

Before you begin

- Ensure that the primary nodes and the cluster are healthy.

To delete an existing secondary node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. Select the checkbox next to the secondary node you want to delete.
4. From the **Actions** menu, choose **Delete** to delete the node.

Managing Standby Nodes

You can add up to two standby nodes, which you can use to quickly restore the cluster functionality in case one or more primary nodes fail by replacing the failed primary node with the standby node.

Standby nodes are similar to secondary nodes in deployment, initial configuration, and upgrades. However, unlike secondary nodes, the cluster will not use the standby nodes for any workloads.



Standby nodes are not supported for single-node clusters or clusters deployed in AWS or Azure.

The following two cases are supported:

- Single primary node failure

You can use the UI to convert the standby node into a new primary node.

- Two primary nodes failure

You will need to perform manual failover of one of the nodes to restore cluster functionality. Then fail over the second node using standard procedure.

Adding Standby Nodes

This section describes how to add a standby node to your cluster for easy cluster recover in case of a primary node failure.

Before you begin

- Ensure that the existing primary nodes and the cluster are healthy.
- Prepare and deploy the new node.

You can failover only between nodes of identical types (physical or virtual), so you must deploy the same type of node as the nodes in your cluster which you may need to replace. In case of virtual nodes deployed in VMware ESX, which have two node profiles (**OVA-app** and **OVA-data**), you can failover only between nodes of the same profile.

- Ensure that the node you are adding is powered on.
- If you are adding a physical node, ensure that you have the new node's CIMC IP address and login information.

You will need to use the CIMC information to add the new node using the Nexus Dashboard GUI.

- If you are adding a virtual node, ensure that you have the node's management IP address and login information.

To add a standby node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. In the main pane, click **Add Node**.

The **Add Node** screen opens.

4. In the **Add Node** screen, provide the node information.

- a. Provide the **Name** of the node.
- b. From the **Type** dropdown, select **Standby**.
- c. Provide the **Credentials** information for the node, then click **Verify**.

For physical nodes, this is the IP address, username, and password of the server's CIMC. The CIMC will be used to configure the rest of the information on the node.

For virtual nodes, this is the IP address and **rescue-user** password you defined for the node when deploying it.

d. Provide the **Management Network** information.

For virtual nodes, the management network information will be pre-populated with the information pulled from the node based on the IP address and credentials you provided in the previous sub-step.

For physical nodes, you must provide the management network IP address, netmask, and gateway now.

e. Provide the **Data Network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

f. (Optional) Provide IPv6 information for the management and data networks.

Starting with release 2.1.1, Nexus Dashboard supports dual stack IPv4/IPv6 for the management and data networks.

If you want to provide IPv6 information, you must do it when adding the node.

All nodes in the cluster must be configured with either only IPv4 or dual IPv4/IPv6 stack.

5. Click **Save** to add the node.

The configuration will be pushed to the node and the node will be added to the list in the GUI.

Replacing Single Primary Node with Standby Node

This section describes failover using a pre-configured **standby** node. If your cluster does not have a standby node, follow the steps described in one of the sections in [Troubleshooting](#) instead.

Before you begin

- Ensure that at least 2 primary nodes are healthy.
- Ensure that you have at least one **standby** node available in the cluster.

Setting up and configuring **standby** nodes is described in [Adding Standby Nodes](#).

- Ensure that the **primary** node you want to replace is powered off.



You cannot re-add the **primary** node you are replacing back to the cluster after the failover is complete. If the **primary** node you replace is still functional and you want to re-add it to the cluster after the failover, you must factory reset or re-image it as described in [Troubleshooting](#) and add it as a **standby** or **primary** node only.

To failover a single primary node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. Click the **Actions (...)** menu next to the **Inactive** primary node that you want to replace.
4. Choose **Failover**.

Note that you must have a standby node already configured and added or the **Failover** menu option will not be available.

5. In the **Fail Over** window that opens, select a standby node from the dropdown.
6. Click **Save** to complete the failover.

The failed primary node will be removed from the list and replaced by the standby node you selected. The status will remain **Inactive** while the services are being restored to the new primary node.

It can take up to 10 minutes for all services to be restored, at which point the new primary node's status will change to **Active**.

Replacing Two Primary Nodes with Standby Nodes

This section describes failover using a pre-configured **standby** node. If your cluster does not have a standby node, follow the steps described in one of the sections in [Troubleshooting](#) instead.

If only one of your primary nodes failed, you can use the GUI to replace it with a standby node as described in [Replacing Single Primary Node with Standby Node](#).

However, when two primary nodes are unavailable, the cluster goes offline. In this case, most operations including the UI are disabled and no configuration changes can be made to the cluster. You can still SSH into the remaining primary node as the **rescue-user**, which is used to recover the cluster by manually failing over one of the failed primary nodes to a standby node. Once two **primary** nodes are available again, the cluster can resume normal operation, at which point you can recover the second primary node using the normal procedure.

One example of when this scenario can occur is in the following topology when nodes are distributed across different sites and a site with 2 **primary** nodes goes down leaving only a single **primary** node and a single **standby** node in the cluster:




Figure 3. 2-node DR Topology

Before you begin

- Ensure that you have at least one **standby** node available in the cluster.

Setting up and configuring **standby** nodes is described in [Adding Standby Nodes](#).

- Ensure that the **primary** nodes you want to replace are powered off.



You cannot re-add the **primary** node you are replacing back to the cluster after the failover is complete. If the **primary** nodes you replace are still functional and you want to re-add them to the cluster after the failover, you must factory reset or re-image them first as described in [Troubleshooting](#) and then add them as a **standby** or **primary** node only.

- If you had installed the Nexus Dashboard Fabric Controller (NDFC) service in the cluster, you must have a configuration backup available to restore after you recover the cluster.

The Fabric Controller service cannot recover from a two **primary** node failure of the Nexus Dashboard cluster where it is running. After your recover the cluster, you must re-install the NDFC service and restore its configuration from a backup.

To fail over two primary nodes:

1. Log in to the remaining primary node via CLI as **rescue-user**.
2. Use the **acs show nodes** command to check the current state of the nodes in the cluster.

This should show two nodes (**M1** and **M2**) as **Inactive** and two other nodes (**M3** and **S1**) as **Active**.

In addition, if an entire site is down and it had any **secondary** or **standby** nodes, they will also

show as **Inactive**.



Ensure that you bring down or disconnect all **Inactive** nodes so they cannot re-establish connectivity to the remaining cluster.

3. Execute the failover command.

In the following command, replace **<M1-node-name>** and **<S1-node-name>** with the names of a failed node and the standby node:

```
# acs failover --failedNode <M1-node-name> --standbyNode <S1-node-name>
```

4. Confirm that you want to proceed.

Warning: Failover can be a disruptive operation and should only be performed as a last resort option to recover cluster from disasters using standby where two primary nodes have lost their state due to hardware faults.

Proceed? (y/n): y

The primary node will copy the configuration state to the standby node and both nodes will restart. It may take up to 30 minutes for the nodes to come up and the cluster to be restored. You can check the progress by navigating to the primary node's UI.

5. Wait for the failover to complete.

The healthy primary node will copy the configuration state to the standby node and both nodes will restart. It may take up to 30 minutes for the nodes to come up and the cluster to become functional. You can check the progress by navigating to one of these node's UI.

6. Remove any inactive **Secondary** or **Standby** nodes from the cluster.

If there are any inactive **Secondary** or **Standby** nodes (for example in the failed site), you must remove them from the cluster.

- a. Navigate to **Admin Console** and choose **Manage > Nodes**.
- b. Click on the three dots next to the **Inactive** nodes choose **Delete**.



If at any point the **secondary** nodes recover and reconnect (such as if the site connectivity is re-established), the nodes will get discovered by the cluster and show in the UI. In that case, you must delete them in the UI again.

7. Ensure the cluster has no errors except ones that are specific to the NDFC service.

Use the **acs recover health-check** command to ensure that the cluster is healthy with only NDFC-specific issues that may still be listed:

```
acs recover health-check
```

Output of this command should look like below before you proceed to Step 5:

```
Health checks in app scope failed. Faults:
[minor] cisco-ndfc-controller-elasticsearch: could not fetch component status
```

8. If you have NDFC service installed, clean wipe NDFC service and restore configuration.

If you don't have NDFC service running, skip this step.

After the previous step is complete, you have an operational two-node ND cluster running with (M3, S1) as active primary nodes. You can now log in to the cluster GUI using either of these two active nodes and restore NDFC. .. Log in to the GUI and open **Admin Console**. .. Navigate to **Analyze > Service Status**. .. Click on the three dots next to Nexus Dashboard Fabric Controller and choose **Clean Wipe**.

+ This will re-install NDFC and bring it up in a clean state. Wait for NDFC to reach **Enabled** state.

- a. Restore NDFC configuration from a previously taken backup.

At this point you have an operational cluster with two master nodes and NDFC service enabled. While we recommend that you add the 3rd primary node as soon as possible to restore cluster redundancy, you can choose to restore NDFC service configuration at this time.

- b. Validate that NDFC is up and running before moving to the next step of re-adding the failed nodes back to the cluster.

You can use the `acs recover health-check` command to ensure that the cluster has no more warnings:

```
acs recover health-check
Health checks in app scope successful
```

9. If this cluster is part of a federation of clusters, remove it from the federation.
 - a. Login to the **Admin Console** on the Primary cluster in the federation.
 - b. Navigate to **Admin > System Settings > Multi cluster connectivity**.
 - c. Remove this cluster from the federation.



If this cluster is the primary cluster of the federation, then navigate to the Multi cluster connectivity and remove all other clusters one at a time.

10. If the failed nodes are still operational and you want to re-add them to the cluster, factory reset them.

If the nodes are not operational, RMA the hardware or bring up another virtual node.

If the nodes are still operational, they can be factory reset by running `acs reboot factory-reset`. The nodes will reboot and lose all configurations and get reset to factory provisioned state.



If you had any secondary nodes that were removed in a previous step, follow the factory-reset steps for them as well.

- o In case these are physical nodes, they will come to a prompt in CIMC host console that says, "Press any key to run first-boot setup on the console". Validate that (M1, M2) management and data IP addresses are not pingable from (M3, S1).
- o In case these are virtual nodes provisioned using VMware vCenter, they will eventually display "System UI online, please login to <node-ip> to continue" on the console. Validate that (M1, M2) data IP addresses are not reachable from (M3, S1). Management IPs should be reachable.

In case these are virtual nodes provisioned directly on an ESXi without using vCenter they will eventually display "Press any key to run first-boot setup on the console" on the console. You should configure the management IP and password by pressing enter here. Validate that (M1, M2) data IP addresses are not reachable from (M3, S1). Management IPs should be reachable.

11. In ND GUI, under the **Admin Console**, navigate to **Manage > Nodes** and note the listed suspended primary node.

You will see only one of the failed master nodes (M1, M2) is still shown and it will be in "Suspended" state. Note the name of this node as this is the first node we will re-add back to cluster. .. On one of the active nodes, navigate to **Manage > Nodes**, click on three dots next to the suspended node and choose **Register**.

+ It may take several minutes for M2 to finish completely joining the cluster. After, the cluster will have 3 primary nodes (M2, M3, S1) and become a fully functional ND cluster.

- a. After the node is registered, verify that the cluster is healthy and the services are still functional.

You can use the **acs health** command on any node to verify that cluster is healthy before proceeding.

12. Perform post-recovery command to ensure ND cluster and services are fully functional.

- a. SSH into one of the primary nodes as **rescue-user** and run **acs recover** command.

This command internally resets the NDFC minio cluster and recovers certain keys on the ND to mitigate GUI login issues.

- b. Confirm that you want to proceed.

Warning: This command executes post disaster recovery actions.
 This command should be run only after:

1. You have successfully recovered two master nodes failures either using "acs rma" or "acs failover" command.
2. You have successfully completed RMA/Register of third master node.
3. You have executed "acs health" and cluster is healthy.Proceed? (y/n): y

.. Wait for the process to complete.

This may take up to 20 minutes.

13. Disable/Enable NDFC from the ND GUI, so that NDFC recovery is complete.

If you do not have NDFC service, skip this step. .. In the **Admin Console**, navigate to **Analyze >**

Service Status. .. Click on the three dots next to Nexus Dashboard Fabric Controller and choose **Disable**. .. Wait for NDFC to reach “Disabled” state .. Click on the three dots next to Nexus Dashboard Fabric Controller and choose **Enable**. .. Wait for NDFC to reach “Enabled” state

+ NOTE: If any image management policies were created in NDFC after you clean wiped the service, images uploaded for the policies must be re-uploaded from the NDFC’s **Image Management** GUI. Policies will reconcile with the newly uploaded images and do not have to be re-created.

14. Re-add a standby node to the cluster.

After the cluster has been restored to a full 3-node cluster as described in the previous step, you can add the last node (M1) as a standby node back to the cluster. .. In **Admin Console**, navigate to **Manage > Nodes** and click **Add Node**.

+ It will take a few minutes for M1 to join the cluster and be healthy.

+ In case you had secondary nodes that were removed in a previous step, ensure they are factory reset and add them back to the cluster.

a. If this cluster was part of a federation of Clusters, re-add it to the federation.

Deleting Standby Nodes

Before you begin

- Ensure that the primary nodes and the cluster are healthy.

To delete an existing standby node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. Select the checkbox next to the standby node you want to delete.
4. From the **Actions** menu, choose **Delete** to delete the node.

Trademarks

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2017–2024 Cisco Systems, Inc. All rights reserved.

First Published: 2024-03-01

Last Modified: 2024-03-01

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<http://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883